

#130
1-30-04

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE
Art Unit 2121
Examiner Joseph P. Hirl

In Re: Rodolfo Milito et al.
Case: P3807
Serial No.: 09/586,115
Filed: June 2, 2000
Subject: Wire-Speed Multi-Dimensional Packet Classifier

To the Commissioner of Patents and Trademarks
Washington, D.C. 20231

RECEIVED

JAN 27 2004

Dear Sir,

Technology Center 2100

Response D

All of the claims standing for examination are reproduced below. There are no amendments to the claims or specification herein made.

1. (Previously Amended) A system for classifying packets, wherein each packet has N header fields to be used for processing, the system comprising:

a first set of rules associating to the packets by values of the header fields;

and

a classification system for selecting specific rules in the set of rules as applicable to a specific packet;

characterized in that the classification system projects the first set of rules as N-dimensional entities on N axes in N-dimensional space, marking the beginning and ending value on each axis for each rule as a breakpoint, assigns a sequence of binary numbers to each interval between breakpoints such that all adjacent intervals are numbered in ascending sequential order, associates a subset

of the first set of rules applicable in each interval to the assigned binary number of the appropriate interval between breakpoints on each axis, then considers a packet as a point in the N-dimensional space according to its header field values, locates the binary number assigned to the interval into which the point projects on each axis by performing a search on each axis for the numbered interval into which the point projects on that axis, thereby determining rules applicable to the packet for that axis, and then determines the specific rules applicable to the packet from the subsets of rules by selecting those rules as applicable to the packet that apply to the packet on all of the N axes.

D. 1
Contd.

2. (Original) The system of claim 1 wherein the search performed on each axis is a binary search conducted by selecting breakpoints at which the bits change for the binary numbered intervals.

3. (Original) The system of claim 1 wherein the search performed on each axis is a quaternary or higher-level M-ary search, where M is a power of 2, conducted by selecting breakpoints at which the bits change for the binary numbered intervals.

4. (Original) The system of claim 1 wherein association of applicable rules in each numbered interval is made by associating a binary string with each interval, with one bit dedicated to each rule.

5. (Original) The system of claim 4 wherein the rules are associated to bit positions in the binary string by priority, the order of priority according to bit significance, and a final rule is selected by the most significant 1 in the matching rules.

6. (Original) The system of claim 4 wherein the applicable rules are found by

ANDing the binary strings determined for each axis over all axes.

7. (Original) The system of claim 1 comprising at least one hardware pipeline for conducting the search on an axis, the pipeline comprising first, second, and sequential modules for accomplishing increasingly particular portions of the search, wherein, after the first module of the sequential modules is used, determined values from the first module pass to the second module, and values for a second packet enter the pipeline at the first module, the pipeline operations proceeding thus sequentially.

11
Cmtd.

8. (Original) The system of claim 7 comprising parallel pipelines with one pipeline dedicated to searching on each axis in the N-dimensional space, wherein searches are conducted for applicable intervals simultaneously on each axis.

9. (Original) The system of claim 8 wherein applicable rules for each interval on each axis are represented by individual bitmaps, with each rule assigned a bit position, and wherein the outputs of the parallel pipelines, being the numbered interval on each axis into which the point for a packet projects, are exchanged for the associated bitmaps, which are then ANDed to determine the applicable rules.

10. (Original) The system of claim 1 wherein searching is interleaved, results of searching on one or more axes being applied to other axes before searching on the other axes.

11. (Original) The system of claim 10 wherein rules that are found by search to not apply on one or more axes are not considered in searches conducted on the other axes.

12. (Previously Amended) A method for classifying packets in routing, wherein each packet has N fields to be used in processing in a header, comprising the steps of:

- (a) projecting the rules as N-dimensional entities on N axes in N-dimensional space;
- (b) marking the beginning and ending value on each axis for each rule as a breakpoint;
- (c) assigning a sequence of binary numbers to intervals between breakpoints on each axis such that all adjacent intervals are numbered sequentially in ascending order;
- (d) identifying those breakpoints at which bits in the interval numbers change;
- (e) associating a subset of the rules as applicable in to the assigned number of each interval on each axis;
- (f) considering a packet as a point in the N-dimensional space according to values of the header fields for the packet;
- (g) determining by search the binary number of the interval on each axis into which the packet point projects;
- (h) substituting the subset of rules that apply for each determined interval; and
- (i) selecting those rules as applicable to the packet that associate to the packet on all of the N axes.

13. (Original) The method of claim 10 wherein, in step (g) the determination is made by a binary search.

14. (Original) The method of claim 12 wherein, in step (g), the determination is made by a quaternary or higher-level M-ary search.

15. (Original) The method of claim 12 wherein, in step (e), association of applicable rules in each numbered interval is made by associating a binary string with each interval, with one bit dedicated to each rule.

16. (Original) The method of claim 15 wherein the rules are mapped to bit positions in the binary string by priority, the order of priority according to bit significance, and a final rule is selected by the most significant 1 in the matching rules.

17. (Original) The method of claim 15 wherein, in step (i), the matching rules are found by ANDing the binary strings determined for each axis over all axes.

18. (Original) The method of claim 12 wherein, in step (g), the search is conducted by sequential modules in at least one hardware pipeline, the pipeline comprising first, second, and sequential modules for accomplishing increasingly particular portions of the search, and wherein, after the first module of the sequential modules is used, determined values from the first module pass to the second module, and values for a second packet enter the pipeline at the first module, the pipeline operations proceeding thus sequentially.

19. (Original) The method of claim 17 comprising parallel pipelines with one pipeline dedicated to searching on each axis in the N-dimensional space, wherein searches are conducted for applicable interval simultaneously on each axis.

20. (Original) The method of claim 19 wherein applicable rules for each interval on each axis are represented by individual bitmaps, with each rule assigned a bit position, and wherein the outputs of the parallel pipeline, being the numbered

interval on each axis into which the point for a packet projects, are exchanged for the associated bitmaps, which are then ANDed to determine the second set of matching rules.

21. (Original) The method of claim 12 wherein, in step (g) searching is interleaved, results of searching on one or more axes being applied to other axes before searching on the other axes.

22. (Original) The method of claim 21 wherein rules that are found by search to not apply on one or more axes are not considered in searches conducted on the other axes.

23. (Original) In a system for classifying packets by binary or higher-level searching for intervals into which rules project on axes, a method for simplifying a search, comprising steps of:

- (a) conducting a first search on one or more axes; and
- (b) using information from the first search to simplify further searching on remaining axes.